# GENERALIZATIONS OF THE BOX-JENKINS AIRLINE MODEL WITH FREQUENCY-SPECIFIC SEASONAL COEFFICIENTS

**John Aston, David Findley, Kellie Wills and Donald Martin, Census Bureau, Washington, DC 20233**

**Key Words: Time Series, Seasonal Adjustment, Forecasting, Filters, Squared Gain**

## Abstract

The Box-Jenkins "airline" model is the most widely used ARIMA model for seasonal time series. Findley, Martin and Wills (2002) previously examined a generalization of the airline model with a more restricted seasonal moving average factor modeling only seasonal effects and with a second-order nonseasonal moving average factor. In this paper, we generalize the seasonal part of the model by associating combinations of the frequencies 1, 2, 3, 4, 5 or 6 cycles per year with individual coefficients. We also consider properties of model-based seasonal adjustment filters obtained from the new models.

## 1. Introduction

Box and Jenkins (1976) developed a two-coefficient time series model, now known as the airline model, which is by far the most widely used ARIMA model for monthly and quarterly macroeconomic time series. The Box-Jenkins airline model for a seasonal time series $Z_t$ with $s \geq 2$ observations per year has the form

$$(1-B)(1-B^s)Z_t = (1-\theta B)(1-\Theta B^s)\varepsilon_t \tag{1}$$

When $\Theta \geq 0$, the airline model can be written

$$(1-B)^2(\sum_{j=0}^{s-1} B^j)Z_t = (1-\theta B)(1-\Theta^{1/s}B)(\sum_{j=0}^{s-1}\Theta^{j/s}B^j)\varepsilon_t \tag{2}$$

Findley, Martin and Wills (2002) substituted a general MA(2) polynomial for $(1-\theta B)(1-\Theta^{1/s}B)$ in (2), yielding their *generalized airline model*,

$$(1-B)^2(\sum_{j=0}^{s-1} B^j)Z_t = (1-aB-bB^2)(\sum_{j=0}^{s-1}c^jB^j)\varepsilon_t \tag{3}$$

---

This paper reports the results of research and analysis undertaken by Census Bureau staff. It has undergone a Census Bureau review more limited in scope than that given to official Census Bureau publications. This report is released to inform interested parties of ongoing research and to encourage discussion of work in progress.

In this model the seasonal sum polynomial has a third coefficient c distinct from the coefficients associated with the other factors in the model.

In the present paper, we investigate airline model generalizations (*frequency-specific* models) that assign separate seasonal coefficients to different seasonal frequencies.

For monthly data, i.e. $s=12$, the model (3) can be generalized by factoring $\sum_{j=0}^{11} c^jB^j$ in terms of frequencies of 1, 2, 3, 4, 5 and 6 cycles per year to obtain a general frequency-specific model,

$$(1-B)^2(\sum_{j=0}^{11} B^j)Z_t = \\ (1-aB-bB^2)\left[\left(1+c_6B\right)\prod_{j=1}^{5}\left(1-2c_j\cos(\tfrac{2\pi j}{12})B+c_j^2B^2\right)\right]\varepsilon_t \tag{4}$$

If the six $c_i$'s are distinct, the model has a different seasonal coefficient for each frequency, for a total of eight coefficients. Eight coefficients cannot be estimated reliably from macroeconomic time series of typical lengths. Therefore we combine the frequencies into two groups, with all frequencies in each group having the same coefficient, to reduce the total number of coefficients in the model to four.

These new models cannot be estimated with standard ARIMA modeling software. We performed the estimation in the object-oriented matrix programming environment Ox (Doornik 2001), using the state space functions in the SSFPack library (Koopman, Shephard and Doornik 1999).

We used Akaike's AIC to compare the fit of competing models. The airline model is a special case of each frequency-specific model; thus each comparison to an airline model has a p-value that is approximated by a $\chi^2$ distribution. With $\theta^A$ and $\theta^F$ denoting the parameter vectors of the airline and frequency-specific models, respectively, $L(\hat{\theta}^A)$ and $L(\hat{\theta}^F)$ the respective maximum likelihoods, and $\dim\theta^A$ and $\dim\theta^F$ the respective number of parameters in the models,

$$\Delta AIC \equiv AIC(\hat{\theta}^A) - AIC(\hat{\theta}^F) = \\ -2\{\ln L(\hat{\theta}^A) - \ln L(\hat{\theta}^F)\} - 2(\dim\theta^F - \dim\theta^A) \tag{5}$$

**Table 1. Modeling and spectrum results: eight series having minimum AIC 5-1 models without a unit root.**

| Series | $\Theta$ | $\sqrt[12]{\Theta}$ | $c_1$ | $c_2$ | $c_2$ frequency | peak | $\chi^2$ p-value from eq. (7) |
|---|---|---|---|---|---|---|---|
| U37AVS | 0.6667 | 0.9668 | 0.9718 | 0.8086 | 4 | largest | 0.005 |
| U36CVS | 0.3135 | 0.9079 | 0.9093 | 0.9591 | 6 | largest | 0.007 |
| U34KTI | 0.6681 | 0.9669 | 0.9735 | 0.9371 | 4 | largest | 0.019 |
| X41020 | 0.6351 | 0.9629 | 0.9756 | 0.9036 | 1 | smallest | 0.039 |
| U39BVS | 0.2718 | 0.8971 | 0.8951 | 0.8335 | 4 | largest | 0.065 |
| U33CVS | 0.4832 | 0.9412 | 0.9479 | 0.7370 | 1 | relatively large | 0.093 |
| U33LVS | 0.6786 | 0.9682 | 0.9893 | 0.9589 | 4 | largest | 0.130 |
| M00190 | 0.5358 | 0.9493 | 0.9390 | 0.9752 | 4 | smallest | 0.133 |

**Table 2. Series descriptions.**

| Series | Descriptions |
|---|---|
| M00190 | Imports of wine and related products |
| X41020 | Exports of cookware, cutlery, house and garden ware |
| U33CVS | Construction machinery manufacturing |
| U33LVS | Pump and compressor manufacturing |
| U34KTI | Electromedical, measuring, and control instrument manufacturing |
| U34EVS | Communications equipment manufacturing, defense |
| U36CVS | Heavy duty truck manufacturing |
| U36HVS | Aircraft engine and parts manufacturing, nondefense |
| U37AVS | Household furniture and kitchen cabinet manufacturing |
| U39BVS | Sporting goods, doll, toy and game manufacturing |

When the airline model is correct,

$$-2\{\ln L(\hat{\theta}^A) - \ln L(\hat{\theta}^F)\} \sim \chi^2_{\dim\theta^F - \dim\theta^A} \tag{6}$$

asymptotically; see Taniguchi and Kakigawa (2000, p. 61.) Thus a given value of ΔAIC has an approximate p-value

$$p_{\Delta AIC} = P\{\chi^2_{\dim\theta^F - \dim\theta^A} \geq \Delta AIC + 2(\dim\theta^F - \dim\theta^A)\} \tag{7}$$

By contrast, AIC differences between correct and incorrect frequency-specific models do not have a simple asymptotic distribution, because no model is a special case of another.

We will be simultaneously comparing the airline model to multiple frequency-specific models; section 5 discusses how we used the asymptotic distribution above to set empirical critical values for AIC differences in this situation.

## 2. 5-1 Models

As mentioned above, typical macroeconomic time series are not long enough to reliably estimate a different seasonal coefficient for each frequency. Therefore we combine the frequencies into two groups, with all frequencies in each group having the same coefficient. One way of doing this is to assign the same seasonal coefficient $c_1$ to all frequencies except one, which receives coefficient $c_2$. This yields six possible frequency-specific models, which we call 5-1 models. We fit each of the six 5-1 models to 75 Census Bureau series, consisting of the Foreign Trade series and series from the M3 Survey of Manufacturers' Shipments, Inventories and Orders, for which an airline model had originally been chosen over other standard ARIMA models.[1] From the seven models for each series (including the airline model), the minimum AIC model was determined. AIC preferred a 5-1 model over the airline model for 21 series. The largest difference from the airline model's

[1] These are the two major categories of Census Bureau series for which an interesting number of series had a lower AIC for model (3) than for model (1); see Findley, Martin and Wills (2002). For other major categories (Retail Trade, Construction), airline models usually had $\Theta^{1/12}$ very close to 1.

AIC was –6.81, associated with a p-value of 0.0045 from Equation 7. This occurred for the M3 series U37AVS for the model with frequency 4 cycles/year associated with coefficient $c_2$. (See Table 2 for series descriptions.)

The seasonal adjustments of series U37AVS from the ARIMA model-based (AMB) decompositions (Hillmer and Tiao, 1982) of the preferred 5-1 model and the airline model are shown in Figure 1. The squared gains of the end-point (concurrent) and mid-point (symmetric) seasonal adjustment filters are shown in Figures 2 and 3 respectively. The squared gains of the airline and 5-1 models are similar except in a broad interval around the frequency 4 cycles/year, where the smaller squared gains of the 5-1 model indicate that its filters are suppressing more frequency components, which results in the smoother seasonal adjustment. For more information about the squared gains, see Findley and Martin (2003).

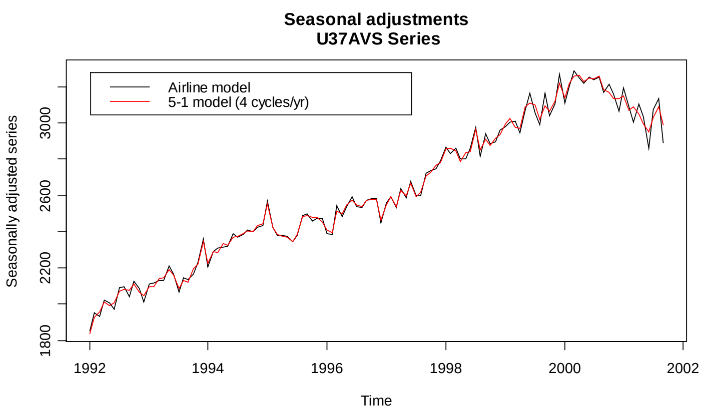


**Figure 1. Seasonal adjustments for series U37AVS. The 4-coefficient frequency-specific model has frequency 4 cycles/year associated with the $c_2$ coefficient.**

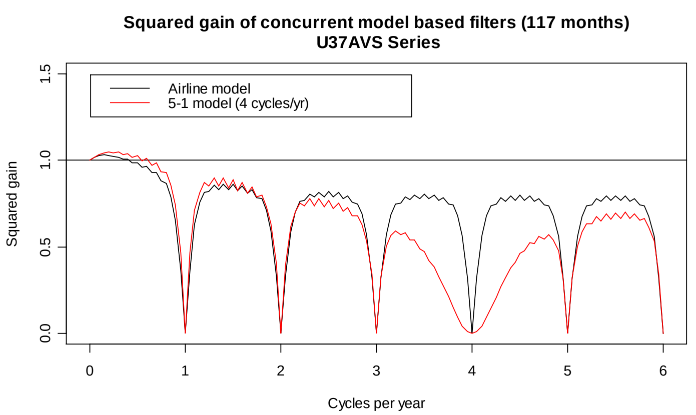


**Figure 2. Squared gain of the finite concurrent model-based filter for series U37AVS.**

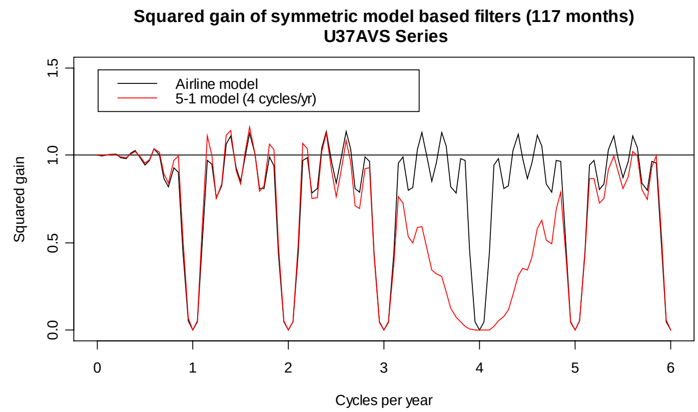


**Figure 3. Squared gain of the finite symmetric model-based filter for series U37AVS.**

For 13 of the 21 series for which AIC preferred a 5-1 model over the airline model, the preferred 5-1 model had one or more unit roots in the seasonal MA polynomial, i.e. either $c_1 = 1$ or $c_2 = 1$. Such unit roots cancel with unit roots in the differencing polynomial, introducing fixed seasonal means and changing the structure of the underlying model so that it is no longer a generalization of the airline model -- its differencing is no longer the same. This changes the structure of the components of the AMB decomposition. We shall defer consideration of models with a unit root for future study, noting for now that the unit roots are often spurious, as we discuss below. Coefficients and other information for the remaining eight minimum AIC 5-1 models without a unit root appear in Table 1.

Seasonal spectral peaks of the (log-transformed and differenced) original series were compared with the frequency associated with the $c_2$ coefficient in the minimum AIC model. In all but a few cases, the frequency associated with $c_2$ in the minimum AIC model corresponds to either the largest or smallest spectral peak. This is graphical confirmation that the $c_2$ frequency in a minimum AIC 5-1 model differs in a consistent way from the other five frequencies. For series U37AVS shown in Figure 1, frequency 4 cycles/year, associated with the $c_2$ coefficient in the minimum AIC model, has the largest peak in the spectrum of the original series. The $c_1$ coefficient estimate is 0.9718, close to the twelfth root of the airline model seasonal coefficient $\Theta$ ($\sqrt[12]{0.6667} = 0.9668$). However, the $c_2$ coefficient estimate is 0.8086. This smaller value results in the squared gains having wider "troughs" at the 4 cycles/year frequency and indicates more variability in the seasonal component around this frequency.

## 3. 4-2 Models

Another four-coefficient version of the frequency-specific model (5) associates four frequencies with $c_1$ and two with $c_2$. There are 15 such 4-2 models. For 37 of the 75 Census Bureau series, AIC preferred at least one of the 4-2 models over the airline model. For 22 of these 37 series, all preferred 4-2 models had unit roots. Among all the preferred 4-2 models that did not have unit roots, AIC differences ranged from –6.77 to –0.49, associated with p-values from 0.005 to 0.106 from Equation 7. Figure 4 shows the seasonal adjustments and Figures 5 and 6 the squared gains of the filters for series U36HVS, with frequencies 1 cycle/year and 3 cycles/year associated with the $c_2$ coefficient. The squared gains of the airline and 5-1 models are similar except in intervals around the frequencies 1 cycle/year and 3 cycles/year, where the larger squared gains of the 4-2 model indicate that its filters are suppressing fewer frequency components. This results in the less smooth seasonal adjustment (particularly visible in the last two years of the series).

Among the preferred 4-2 models, there was rarely an obvious relationship between the size of the peaks in the spectrum of the original series and the pair of frequencies associated with the $c_2$ coefficient. Thus the spectrum does not confirm that the $c_2$ frequencies in the preferred 4-2 models differ from the other four frequencies in a consistent way.

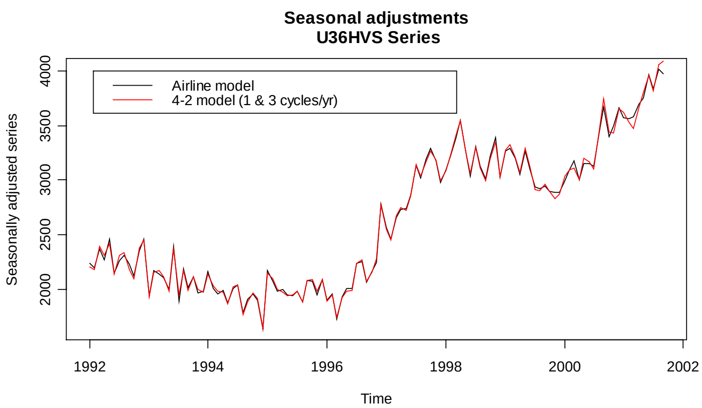


**Figure 4. Seasonal adjustments for series U36HVS. The 4-coefficient frequency-specific model has frequencies 1 cycle/year and 3 cycles/year associated with the $c_2$ coefficient.**

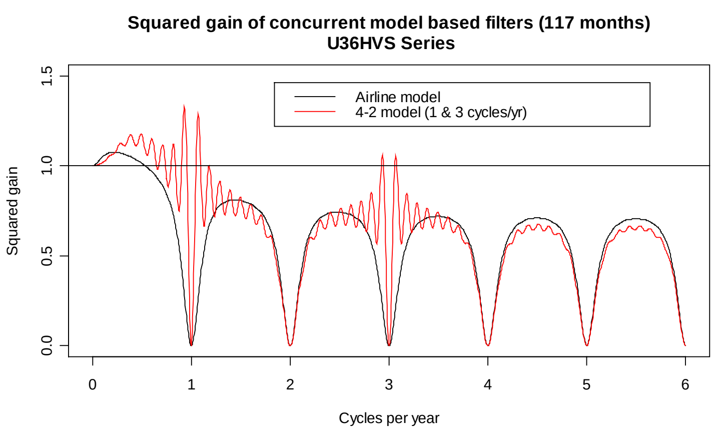


**Figure 5. Squared gain of the finite concurrent model-based filter for series U36HVS.**

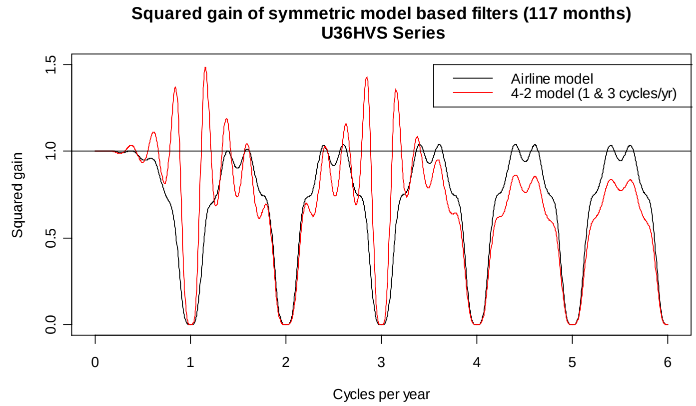


**Figure 6. Squared gain of the finite concurrent model-based filter for series U36HVS.**

## 4. Coefficient Estimation Uncertainties and a Three-Coefficient 5-1 Model

Cryer and Ledolter (1981) first showed that maximum likelihood coefficient estimates of an *invertible* MA(1) model take on a unit magnitude value with a positive probability, i.e. the estimated model is *noninvertible*. Tanaka (1996) provides a detailed theory for this phenomenon. His Table 8.2 (p. 313) presents these probabilities for several sample sizes, as well as for seasonal MA(1) models with $s = 4,12$. Many unit root estimates might be instances of this phenomenon. We performed simulation experiments to determine how accurately the $c_2$ coefficient can be estimated, with particular attention to how often it is estimated as unity, since unit roots change the underlying model, as discussed above.

We generated 1000 simulations of the 5-1 model with series length 150 observations and true coefficients a = 0.5, b = 0.5, $c_1$ = 0.96 and $c_2$ = 0.93. (These are the means of the coefficient estimates for $c_1$ and $c_2$ for the 21 Census Bureau series for which AIC preferred a 5-1 model to the airline model.) Figures 7 and 8 are histograms of the $c_1$ and $c_2$ estimates for the 5-1 model. While the $c_1$ estimates have a fairly narrow distribution around the true value of 0.96 (Figure 7), the $c_2$ estimates are much more widely spread around the true value of 0.93 (Figure 8). In particular, $c_2$ was estimated as unity for 4.7% of the series.

We performed another simulation to investigate whether associating two frequencies with the $c_2$ coefficient (the 4-2 model) resulted in more accurate estimation of $c_2$ (Figure 9). The true coefficients were a = 0.5, b = 0.5, $c_1$ = 0.96 and $c_2$ = 0.93. The 4-2 model reduced the number of estimates in the tails of the distribution, including the number of estimates at unity.

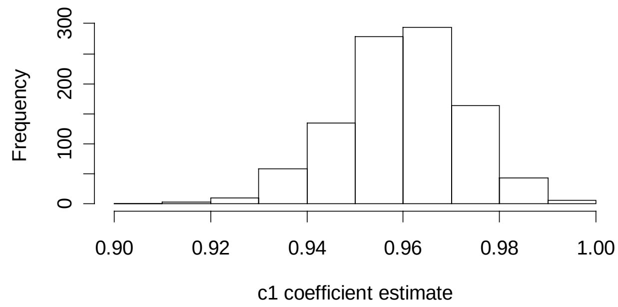


**Figure 7. Distribution of estimates of $c_1$ for 1000 realizations of the 5-1 model with coefficients a = 0.5, b = 0.5, $c_1$ = 0.96, $c_2$ = 0.93. Seven coefficient estimates were unity. Two coefficient estimates were less than 0.9.**

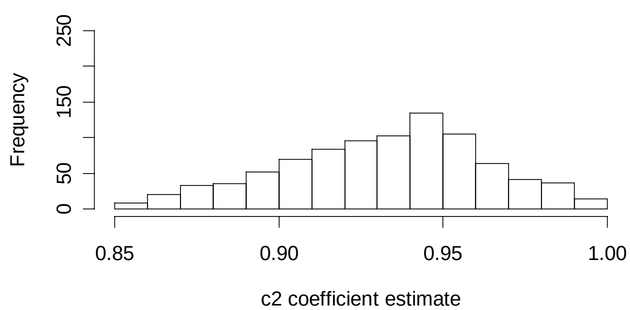


**Figure 8. Distribution of estimates of $c_2$ for 1000 realizations of the 5-1 model with coefficients a = 0.5, b = 0.5, $c_1$ = 0.96, $c_2$ = 0.93. Forty-seven coefficient estimates were unity. Fifty-four coefficient estimates were less than 0.85.**

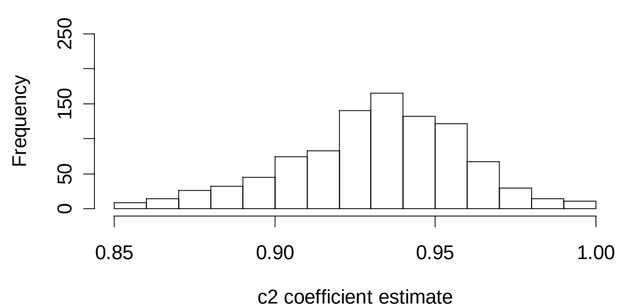


**Figure 9. Distribution of estimates of $c_2$ for 1000 realizations of the 4-2 model with coefficients a = 0.5, b = 0.5, $c_1$ = 0.96, $c_2$ = 0.93. Eleven coefficient estimates were unity. Twenty-four coefficient estimates were less than 0.85.**

Many Census Bureau time series might not be long enough to allow accurate estimation of the four coefficients of the 5-1 model (the Foreign Trade series are 155 observations long, and the M3 series are 117 observations). Thus we developed a three-coefficient version of the 5-1 model in which a root reciprocal of the nonseasonal MA(2) polynomial is forced to have the coefficient $c_1$ associated with five of the seasonal frequencies. (In this version, the MA(2) is restricted to have real roots.) For example, when $c_2$ is assigned to the seasonal frequency 6 cycles/year, the model is

$$(1-B)^2\left(\sum_{j=0}^{11} B^j\right)Z_t = (1-aB)(1-c_1B) \times\left\{(1+c_2B)\prod_{j=1}^{5}\left(1-2c_1\cos\left(\tfrac{2\pi j}{12}\right)B+c_1^2B^2\right)\right\}\varepsilon_t \qquad (8)$$

We performed a simulation experiment to investigate the probability of estimating a unit root with the new parameterization. We generated 1000 realizations with true coefficient values a = 0.5, $c_1$ = 0.96 and $c_2$ = 0.93 (for comparison to the earlier simulations). A histogram of the coefficient estimates appears as Figure 11. For 4% of the realizations, the $c_2$ coefficient was estimated as unity, not a great improvement over the 4.7% estimated as unity for the four-coefficient 5-1 model.

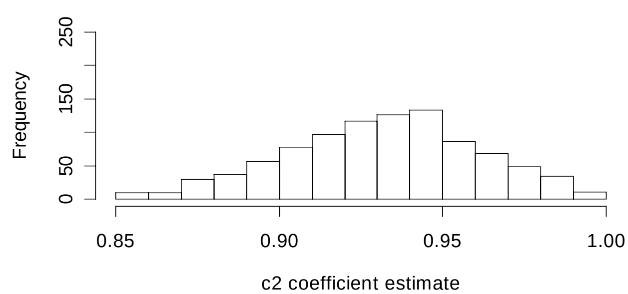


**Figure 11. Distribution of estimates of $c_2$ for 1000 realizations of the 3-coefficient 5-1 model with coefficients a = 0.5, $c_1$ = 0.96, $c_2$ = 0.93. 40 coefficient estimates were unity. 16 coefficient estimates were less than 0.85.**

There are six possible three-coefficient models assigning coefficient $c_2$ to a single frequency. We fit each of these models to the 75 Census Bureau series previously mentioned. For 33 series, AIC preferred at least one of the three-coefficient frequency-specific models to the airline model. AIC differences ranged from –8.68 to –0.52, associated with p-values from 0.001 to 0.113 from Equation 7. Nine of these models had unit roots. Seasonal adjustments and squared gains for series U34EVS, with frequency 4 cycles/year associated with parameter $c_2$, are shown in Figures 12, 13 and 14. The squared gains of the airline and three-coefficient 5-1 models are similar except in a broad interval around the frequency 4 cycles/year, where the smaller squared gains of the three-coefficient 5-1 model indicate that its filters are suppressing more frequency components.

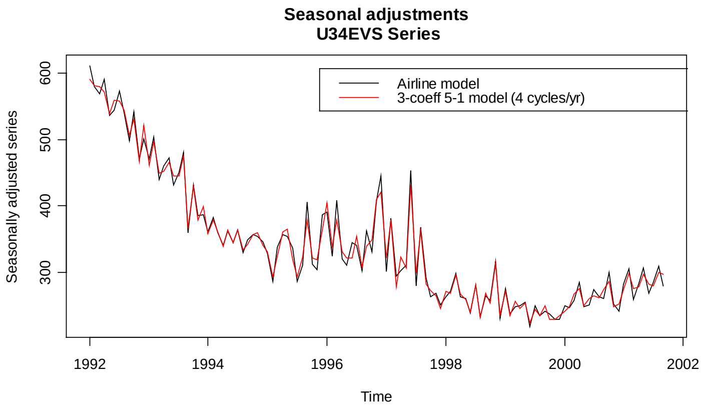


**Figure 12. Seasonal adjustments for series U34EVS. The 3-coefficient frequency-specific model has frequency 4 cycles/year associated with the $c_2$ coefficient.**

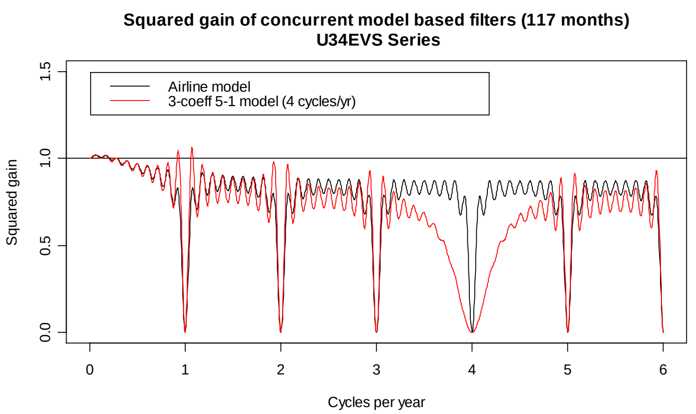


**Figure 13. Squared gain of the finite concurrent model-based filter for series U34EVS.**

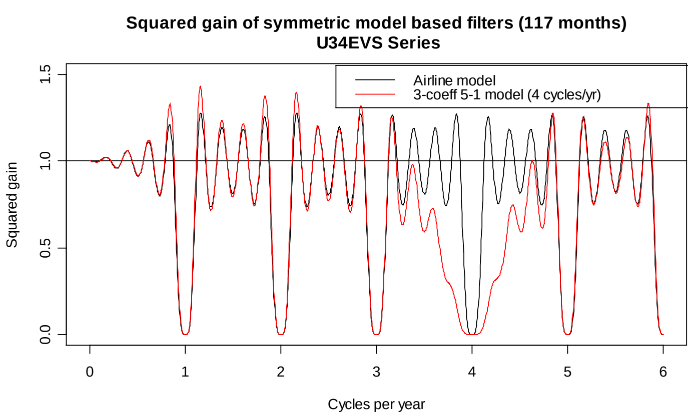


**Figure 14. Squared gain of the finite symmetric model-based filter for series U34EVS.**

## 5. AIC Differences and Multiple Models

We have discussed three frequency-specific generalizations of the airline model: a four-coefficient 5-1 model, a four-coefficient 4-2 model and a three-coefficient 5-1 model. For 33 of the 75 Census series, AIC prefers at least one of these three models (without a unit root) over the airline model. We determined the overall minimum AIC model for each of these series. A three-coefficient 5-1 model has the minimum AIC for 19 series. A four-coefficient 4-2 model has the minimum AIC for 11 series, and a four-coefficient 5-1 model has the minimum AIC for the other three series.

The question arises of appropriate critical values for the AIC differences, taking into account the fact that multiple models are being considered. In the case of the 5-1 models, we are comparing 6 different models to the airline model, and in the case of the 4-2 models there are 15 comparisons. From the results in Section 1, if the airline model is the correct model, the probability that the frequency-specific model will have a smaller AIC (and thus be incorrectly preferred) is

$$P\{AIC(\hat{\theta}^A) - AIC(\hat{\theta}^F) > 0\} = \\ P\{\chi^2_{\dim\theta^F - \dim\theta^A} > 2(\dim\theta^F - \dim\theta^A)\} \tag{9}$$

We can use this probability to empirically set critical values for AIC differences for the multiple comparisons. Here the critical values are chosen to provide comparisons of the models in a way as close as possible to AIC for a single model situation. We simulated 500 series of length 150 observations and 500 series of length 120 observations from the airline model. The series lengths were chosen to correspond to the lengths of the Foreign Trade series (144 observations) and M3 series (117 observations). We fit both an airline model and one of the new models to each of the simulated series and observed the distribution of AIC differences. We used $\theta = 0.5$, $\Theta = 0.5$.

In the case of the 4-coefficient models, $\dim\Theta^F - \dim\Theta^A = 2$, so the probability in Equation (9) is $P(\chi_2^2 > 4) = 0.135$. We determined a threshold AIC difference (airline model AIC – new model AIC) which was exceeded by approximately 13.5% of simulated series. In the case of the 3-coefficient models, $\dim\Theta^F - \dim\Theta^A = 1$. $P(\chi_2^2 > 2) = 0.157$, so we determined an AIC difference which was exceeded by approximately 15.7% of simulated series. The threshold AIC differences appear in Table 3, along with the total number of preferred models and the number for which the AIC difference exceeds the appropriate threshold.

Overall, of the 75 Census series, we will treat 14 as being better fit by one or more frequency-specific models than by the airline model, by the criteria of this section. (The last row of Table 3 sums to more than 14 because several series are better fit by more than one frequency-specific model than by the airline model.) It should be noted that while this number takes into account the multiple models that comprise one frequency-specific generalization, it does not account for the multiple comparisons across either the frequency-specific types or the 75 data sets.

**Table 3. Threshold AIC values from Equation (9). "Total preferred" is the number of series for which the model had a lower AIC than the airline model (no unit roots).**

| Model | *4-coefficient* | | *3-coefficient* |
|---|---|---|---|
| | **5-1** | **4-2** | **5-1** |
| Threshold AIC difference (120 obs) | 1.9 | 2.0 | 1.8 |
| Threshold AIC difference (150 obs) | 1.6 | 2.1 | 2.3 |
| Total preferred | 8 | 15 | 24 |
| Number with AIC differences exceeding threshold | 4 | 7 | 9 |

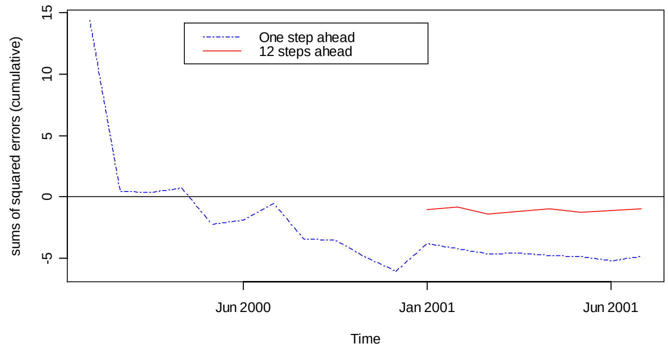


**Figure 15. Out-of-sample forecast error diagnostic for series U34EVS (accumulating sums of squared errors of the airline model subtracted from those of the three-parameter 5-1 model).**

## 6. Forecasting Performance

To obtain information about a model's forecasting performance, some number of observations at the end of the series may be regarded as future data to be forecast from a model fit to the earlier data. These forecasts can be compared to the actual series values, or, for series values identified as outliers, to the outlier-adjusted values. The span of modeled data may be increased one observation at a time, to produce a sequence of h-step-ahead forecast errors. When forecast errors are available from two competing models, the sequence of differences between the accumulating sums of squared errors can be an effective model-selection diagnostic. An example of the diagnostic is Figure 15, showing the sequence of differences between the accumulating sums of h-step-ahead squared errors (h = 1, 12) for series U34EVS. The sums for the airline model are subtracted from those for the three-parameter 5-1 model. The descending dotted line indicates that the one-step-ahead forecast performance of the 3-coefficient 5-1 model is generally better than that of the airline model. The results for 12-step-ahead forecast performance are inconclusive. For more information about the out-of-sample forecast error diagnostic, see Findley, Monsell, Bell, Otto and Chen (1998).

We examined this diagnostic for the models whose AIC difference from the airline model exceeds the appropriate empirically determined threshold discussed in section 5. In many cases, the diagnostic was indeterminate, preferring neither model's forecasting performance over the other. The cases for which the diagnostic preferred one model over the other tended to be rather evenly split between the airline model and the frequency-specific model. However, more 3-coefficient 5-1 models were preferred for lag 1 performance (5 to 1) and more 4-coefficient 4-2 models were preferred for lag 12 performance (3 to 0).

## 7. Conclusions

We have examined frequency-specific airline model generalizations in order to determine what gains may be achieved by modeling one or more seasonal frequencies separately from the others. In order to compare the new models to the airline model, we have examined criteria such as model fit, seasonal adjustments, squared gains and estimation accuracy. By the criteria of Section 5, the new models result in a significant AIC difference for about 19% of the series examined. The one-step-ahead forecast error diagnostic does not suggest strong forecast performance gains for the new models. The new models present certain estimation difficulties: they can't be estimated with standard ARIMA modeling software, and they are prone to problematic unit roots in the seasonal, as well as inaccurate estimation of the $c_2$ parameter. Our experience with these models strengthens our confidence in the robustness and flexibility of the airline model.


## Acknowledgements

The authors are indebted to Kathleen McDonald-Johnson for comments that improved this paper.


## References


Box, G. E. P. and Jenkins, G. M. (1976). *Time Series Analysis: Forecasting and Control* (revised edition), Holden Day, San Francisco.

Cryer, J. D and Ledolter, J. (1981). "Small sample properties of the maximum-likelihood estimator in the first-order moving average model." *Biometrika* **68**, 191-194.

Doornik, J. A. (2001), *Object-Oriented Matrix Programming using Ox 3.0.* London: Timberlake Consultants Press.

Findley, D. F. and Martin, D. E. K. (2003). "Frequency domain analyses of SEATS and X-11/12-ARIMA seasonal adjustment filters for short and moderate-length time series," Research Report S2003-02, Statistical Research Division, U.S. Census Bureau, http://www.census.gov/srd/papers/pdf/rrs2003-02.pdf

Findley, D. F., Martin, D. E. K. and Wills, K. C. (2002). "Generalizations of the Box-Jenkins Airline Model," 2002 Proceedings of the American Statistical Association, Business and Economic Statistics Section [CD-ROM], Alexandria, VA: American Statistical Association.

Findley, D. F., Monsell, B. C., Bell, W. R., Otto, M. C. and Chen, B.-C. (1998), "New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program," *Journal of Business and Economic Statistics*, **16**, 127-152.

Koopman, S. J., N. Shephard, and J. A. Doornik (1999). "Statistical algorithms for models in state space form using SsfPack 2.2." *Econometrics Journal* 2, 113-66.

Tanaka, K. (1996). *Time Series Analysis*, Wiley, New York.

Taniguchi, M. and Kakizawa, Y. *Asymptotic Theory of Statistical Inference for Time Series.* Springer-Verlag, New York.